

The dilemma of Saussurean communication

Michael Oliphant*

Department of Cognitive Science, University of California, San Diego, CA, USA

Abstract

A Saussurean communication system exists when an entire communicating population uses a single 'language' that maps states unambiguously onto symbols and then back into the original states. This paper describes a number of simulations performed with a genetic algorithm to investigate the conditions necessary for such communication systems to evolve. The first simulation shows that Saussurean communication evolves in the simple case where direct selective pressure is placed on individuals to be both good transmitters and good receivers. The second simulation demonstrates that, in the more realistic case where selective pressure is only placed on doing well as a receiver, Saussurean communication fails to evolve. Two methods, inspired by research on the Prisoner's Dilemma, are used to attempt to solve this problem. The third simulation shows that, even in the absence of selective pressure on transmission, Saussurean communication can evolve if individuals interact multiple times with the same communication partner and are given the ability to respond differentially based on past interaction. In the fourth simulation, spatially organized populations are used, and it is shown that this allows Saussurean communication to evolve through kin selection.

Keywords: Saussurean communication; Simulations; Prisoner's Dilemma

1. Saussurean communication

When individuals communicate with one another, the specific symbols (sounds, written characters, etc.) they use are not important. What is important is that each symbol 'means' the same

thing to both the individual sending it and the individual receiving it. It must be possible to map some concept onto a symbol and then map back from the symbol to get the original concept. This bi-directional mapping between meaning and symbol is called a Saussurean sign, after Saussure (1959).

Given that such communication systems seem desirable, it becomes important to understand how they might arise. Because Saussurean com-

* Corresponding author, Fax: +1 619 5341128; E-mail:oliphant@cogsci.ucsd.edu.

munication provides an ideal situation where every individual can perfectly understand every other individual, it seems obvious that populations should develop such systems. It is critical to realize, however, that the fact that a state of affairs is ideal does not explain how it might have evolved. Natural selection occurs at the level of the individual, not at the level of the entire population. It does not directly act to push the population toward some ideal, but rather favors those individuals that do the best they can for themselves. This paper describes investigations using simulations of evolutionary processes to examine the conditions that are necessary for a Saussurean communication system to evolve in populations of self-interested individuals.

2. Simulating evolution

It is very difficult to provide non-speculatory accounts of the evolution of an ability. Computational simulations are useful for attacking this problem. Simulations provide a principled way of testing hypotheses about what pressures may have driven the evolution of a particular phenomenon. Although simulations are necessarily simplified, the hope is that if the simulation captures the important aspects of the phenomenon of interest then the results can provide information about what might have happened in real-world evolution. Such simulations have proved useful in studying many different phenomena, such as the interaction between learning and evolution (Hinton and Nowlan, 1987; Nolfi, Elman, and Parisi, 1990; Parisi, Nolfi, and Cecconi, 1991; Ackley and Littman, 1991) and the impact of culture on these processes (Belew, 1990; Hutchins and Hazelhurst, 1991), the evolution of foraging behavior (Collins and Jefferson, 1991), and the evolution of cooperation (Axelrod, 1984, 1987; Lindgren, 1991).

A number of simulations have been conducted using both learning and evolution to investigate how communication systems might arise. Yanco and Stein (1992) used task-based reinforcement to allow mobile robots to develop a communication protocol to coordinate their activity. Werner

and Dyer (1991) used a genetic algorithm to evolve mating signals that allowed immobile 'females' to direct mobile but blind 'males' to their location. Hurford (1989) studied the evolution of the Saussurean sign and proposed it as a part of the brain mechanism used to learn language. The work described in this paper uses similar techniques to study the conditions required for Saussurean communication to evolve.

3. Simulations

The present simulations use a genetic algorithm (Holland, 1975) to simulate the evolution of communication ability in a population over the course of many generations.

3.1. *Simulation 1: evolving Saussurean communication*

The goal of this set of simulations was to provide a simple example of a population converging to a Saussurean communication system. The task involved communication between two individuals. One individual was designated to be the 'transmitter.' This individual produced a one-bit symbol based on a one-bit environmental state. Thus, the transmitter could send a symbol of '0' or '1' based on an environmental state of '0' or '1'. The other individual, designated to be the 'receiver,' produced a one-bit response to the one-bit symbol provided by the transmitter. Successful communication between the two individuals was considered to have occurred when the receiver's response to a symbol produced by the transmitter matched the environmental state the transmitter was observing.

For each simulation, a random population of 100 individuals was created. The genome for an individual consisted of four bits. These four bits encoded both a transmission strategy (the first two bits) and a reception strategy (the second two bits). The structure of the '1010' genome can be seen in Fig. 1.

Fitness in the simulations was based upon an individual's average communicative success. Both the sender and the receiver were 'rewarded' when communication between them was successful, and

Transmission Genes		Reception Genes	
Symbol		Response	
Env	0	1	Symbol
	0	1	0
	1	0	1
			0

Fig. 1. Structure of the '1010' genome. An individual with this genome has a transmission system ('10') that will transmit the symbol '1' when given the environmental state '0', and the symbol '0' when given the environmental state '1'. The reception system for this genome ('10') produces a response of '1' when given the symbol '0' and '0' when given the symbol '1'. Note that this is a Saussurean communication system, as the reception system is the inverse of the transmission system.

	Match	Mismatch
Transmitter	1	0
Receiver	1	0

Fig. 2. Payoff Matrix for Simulation 1. Both the sender and the receiver get a high (1) payoff when communication between them is successful, and a low ('0') payoff when it is not.

'punished' when it was not. This was done via a very simple payoff matrix, shown in Fig. 2. Based on this payoff scheme, the best fitness an individual could hope to achieve would be 1.0 (resulting from successful communication in every interaction) and the worst fitness would be 0.0 (resulting from communication failure in every interaction).

The evaluation phase of a generation consisted of each member of the current population playing against 16 other randomly-selected members of the population. An interaction consisted of two communication attempts — each individual taking a turn at being the transmitter and a turn at being the receiver. This provided at least 16 evaluations of each genome in each role as transmitter and as receiver, plus an average of 16 more evaluations due to being randomly selected to play with other members of the population.

Each new generation was created from the previous generation by performing 100 random selections biased by fitness (meaning that if one individual had twice the fitness of another, it would be twice as likely to be selected and would be expected to be represented twice as many




Fig. 3. Frequency of the Saussurean communication system '0101' when selective pressure is on both sender and receiver.

times in the new population). The new population was then subjected to mutation (with each bit of each individual's genome having a 27% chance of being mutated) and crossover (with each individual having a 10% chance of being involved in a crossover)¹.

Numerous runs with different initial populations were done. In each case, the entire population quickly converged to a single transmission/reception system. Two such stable states exist — the two Saussurean communication systems possible with this four-bit genome ('0101' and '1010'). Which of them the population converges to depends on the random seed given to the simulator. An example run where the population converged to an '0101' communication system can be seen in Fig. 3. After approximately 40 generations, almost every individual in the population has the same communication system. The variability seen after generation 40 is due to the mutations occurring in each generation.

3.2. Simulation 2: failure to evolve Saussurean communication

In Simulation 1, the consequences of a success-

¹The values for the mutation and crossover rates are arbitrary. A range of values produce comparable results. The key is to have the rates high enough to produce variability in the population while not having them so high that good solutions are quickly lost due to overwhelming variation.

ful communication were the same for both the sender and the receiver. If communication is successful, both benefited. If communication failed, both paid the price. Communication in the real world may not always reflect this assumption. An example is the case of animal alarm calls. Vervet monkeys have a fairly elaborate alarm call system. They have calls that differentially signal the presence of various predators (such as leopards, eagles, and snakes) and produce specific responses in other monkeys that are appropriate to the type of predator (Cheney and Seyfarth, 1990). Vervets certainly seem to have a Saussurean communication system, but it is not clear that the environment in which it evolved is like the simulations just described.

The main problem is that the transmitter of the alarm call appears to get no direct benefit from successful communication — after all, the monkey giving the signal has already seen the predator. The benefit to the receiver is much more clear. If the receiver does not understand the signal, it stands a higher chance of getting eaten. With this in mind, several additional simulations were performed. These new simulations were identical to the earlier ones, except that instead of evaluating both the transmitter and the receiver based on the success of an attempted communication, only the receiver was evaluated. The new payoff matrix is shown in Fig. 4.

After a thousand generations, populations run with this new payoff matrix showed no signs of converging to a single communication system. The data revealed that while the transmission systems showed no tendency to converge, the receiving systems did. Fig. 5 shows that, although the population was almost always converged to one reception system ('01' or '10'), the particular system

	Match	Mismatch
Transmitter	n/a	n/a
Receiver	1	0

Fig. 4. Payoff Matrix for Simulation 2. Selective pressure is placed only on the receiver.




Fig. 5. Frequency of reception systems '01' (upper plot) and '10' (lower plot) when no selective pressure is placed on the transmitter. Note that the population is always dominated by one or the other.

that was dominant was not stable over time². Note the sharp transitions from an '01' reception system to a '10' reception system and back again over the course of the run.

This bistable nature of the reception system distribution reflects the nature of the fitness function. Because individuals are evaluated based on their success at reception, it is advantageous to have a reception system of '01' or '10,' because the other systems give at best chance performance. The reason the population is always converged to one of these systems is because reception is reflecting the transmission system domi-

² Note that the figures are from a particular simulation, but the behavior was not unique to this run.

nant in the population at the time. If one reception system has even a slight advantage over the other, it is quickly converged upon. The other reception systems ('00' and '11') show consistent low frequency.

The transmission systems show no clear domination by any one system. This occurs because there simply is no pressure to force the transmission system to converge. No advantage is given by having a transmission mechanism that corresponds to the reception mechanism of others in the population. Consider two communication systems that both have the same reception system ('01'). Suppose that one of the communication systems has a transmission system that is the 'Saussurean other half' of their common reception system ('01'), while the other communication system has some other reception system (say, '00'). These two communication systems will have identical values for their expected fitness. This is because they will all perform identically in reception, and that is all that that matters to the fitness function in this simulation.

What happens in Simulation 2 seems very similar to the events that unfold in simulations of the Prisoner's Dilemma, where populations of individuals playing one-shot games rapidly converge to defection. In communication, having a transmission system that provides information that is not useful can be seen as a form of defection. Given this, it seems likely that methods that solved the problem of defection might also solve the problem of evolving altruistic communication. One variant typically used in the Prisoner's Dilemma is an iterated game, where individuals play more than one round against an opponent. In order to benefit from multiple interactions with the same opponent, individuals are given a three-round history that documents both their actions and those of their opponent, and are also given a mechanism that allows them to modify their future behavior based on the past behavior of their opponent. As has been shown numerous times, this allows a system of reciprocal altruism to evolve (Axelrod, 1984, 1987, Lindgren, 1991).

Another possibility is that the pressure to be

altruistic operates at the level of the gene, rather than at the level of the organism. This can allow apparently altruistic behavior by an organism to evolve because it is ultimately selfish from the point of view of the organism's genes (Dawkins, 1989). Recent work has shown that organizing a population spatially can cause a 'selfish gene' effect to come into play, and can lead to cooperation in the non-iterated Prisoner's Dilemma (Oliphant, 1994) and more generous cooperative strategies in the iterated game (Grim, this issue). The possibility that both of these methods might also prove effective in the area of communication is explored in the next two sections.

3.3. Simulation 3: reciprocal altruism

The third set of simulations were similar to those without direct selective pressure on the transmitter (Simulation 2), with modifications to allow for an iterated game. Each individual had two four-bit communication systems and a single default memory bit. Instead of playing a single round against an opponent, 16 rounds were played. Each player used one of their two communication systems based on the result of the last interaction (one communication system was used if communication was successful the last time they received their opponent's transmission, and the other was used if communication failed). The memory bit in the genome coded for an initial assumption about the other player and was used to determine which of an individual's two communication systems would be used in the first round that was played against a new opponent (the result of the previous round could not be used, because there was not yet any history of interaction between the two individuals).

Results show that individuals soon began to 'trust' one another, with the population soon converging to the memory bit being '1', indicating an initial assumption that the opponent will be transmitting accurately (see Fig. 6). Soon after this 'trust' develops, one of the Saussurean communication systems ('1010' in this case) quickly comes to dominate the population's primary communication system (see Fig. 7). This is termed the primary communication system because it is the




Fig. 6. Frequency of memory bit '1', which indicates an initial assumption that the opponent will be transmitting accurately in the iterated game.




Fig. 8. Frequency of secondary (selected if communication in the last round was unsuccessful) transmission system '10' in the iterated game.




Fig. 7. Frequency of the Saussurean primary (selected if communication in the last round was successful) communication system '1010' in the iterated game.

communication system that is selected by the default memory bit and it is the system that will be used as long as successful communication continues.

The other half of the strategy that emerges is that the transmission system of the secondary communication system (the one not selected by the memory bit and the one used if an opponent is not 'cooperating') becomes anything but the transmission system that corresponds to the dominant primary reception system of the population (in the example shown, anything but '10' — see Fig. 8). This basically implements an 'I won't give you a good signal if you don't give me one' strategy. Having a 'nice' (Saussurean) primary

communication system and a 'nasty' (non-Saussurean) secondary communication system allows an individual to get a perfect score when playing against others of its kind, while remaining resistant to individuals whose primary communication system is non-Saussurean.

This reciprocally altruistic strategy is not completely stable, however, and in this run seems to fall apart soon after the 400th generation. By generation 100, all individuals in the population are consistently using a Saussurean primary communication system. Because of this, the secondary communication system is no longer used, and hence is no longer subject to selective pressure. As can be seen in Fig. 8, the secondary transmission system begins to drift toward '10,' which is undesirable because it is accurate, causing non-cooperators to go unpunished. This allows non-cooperators to infiltrate the population, and causes a crash in frequency of the '1010' primary communication system and the 'trusting' memory bit. Later in the simulation, it can be seen that a more careful cooperative strategy emerges, and the population again begins to communicate optimally.

3.4 Simulation 4: spatially organized populations

Although the previous set of simulations resulted in the evolution of a Saussurean communication system without direct selective pressure on the transmitter, the actions of individuals were

not truly altruistic. Each individual was selected for their ability to maximize their own fitness; it simply turned out that it was in each individual's best interest to be 'nice.' Using spatially organized populations, on the other hand, has the potential to result in true altruistic behavior at the level of the individual.

In the previous simulations, the populations were not spatially organized. This meant that when picking an opponent to communicate with, or a mate to perform the genetic crossover operator with, all other individuals were equally likely to be chosen. With a spatially organized population, this is not the case. The population can be thought of as existing in a one-dimensional space. Individuals are more likely to communicate and mate with those close to them than they are with those farther away. Also, when an individual has offspring, they are placed in the area of the space where the parent was. These factors result in a space where individuals are more related (genetically closer) to those nearer to them. Because individuals communicate more with those around them, if they are 'nice' (transmit signals that correspond to the dominant reception system in the population), this will benefit those close to them. Since those close to them will also be more related to them, they are, in a sense, benefitting their own genes. The notion that individuals might exhibit behavior that is not in their best interest, but benefits others that are related to them has been a prominent idea in ethology for quite some time (Hamilton, 1964; Dawkins, 1989).

A set of simulations was carried out to test whether such a situation could help lead to the evolution of a Saussurean communication system. These simulations were identical to those in Simulation 2, except that spatially organized populations were used. Spatial organization was imposed by selecting partners to communicate and mate (perform crossover) based on a Gaussian distribution around that individual. In addition, when a new population was created from the previous population, offspring were placed in the same spatial area that their parents occupied. This maintains a spatial organization across generations. Fig. 9 shows the results of a sample




Fig. 9. Frequency of communication system '0101' in a simulation where individuals were biased toward communicating and mating with individuals that were, on average, a distance of five individuals on either side of them.

simulation done where individuals were biased toward communicating and mating with individuals that were, on average, a distance of five individuals on either side of them. After a hundred generations or so, the Saussurean communication system '0101' dominates the population. This result is even more pronounced when a simulation is run with individuals communicating and mating only with those immediately on either side of them, while increasing this distance shows the same general pattern, although it is less stable. In general, the larger the distance, the more the results resemble those obtained with populations that are not spatially organized.

From these simulations, it seems quite clear that organizing a population spatially makes a critical difference. The only difference between these simulations and those in Simulation 2 was the addition of spatial organization. Saussurean communication did not evolve in Simulation 2, but it did evolve in Simulation 4. In Simulation 2 it was shown that there was no advantage to having a 'nice' transmission system — one that corresponds to the reception mechanism of others in the population. This was because an individual's 'niceness' could not differentially benefit 'nice' individuals and 'nasty' individuals. This is no longer the case once spatial organization is imposed. As was noted earlier, 'nice' individuals will end up being close together and will benefit

each other. On the other hand, 'nasty' individuals will also end up being close together and will hurt each other. This leads to a situation where 'nice' individuals flourish and 'nasty' individuals perish, and explains the convergence to a Saussurean communication system.

4. Discussion

The simulations described in this paper have demonstrated that, while a mutually interpretable communication system where everyone understands what everyone else is saying seems obviously desirable, certain conditions must be met in order for such a Saussurean communication system to evolve. In particular, there must be pressure to select for good transmission systems (as was shown by the failure to evolve Saussurean communication in Simulation 2). Because it seems unlikely that such pressure is direct (as was the case in Simulation 1), it must be provided indirectly. Two such indirect methods, both inspired by the Prisoner's Dilemma, were explored: reciprocal altruism (in Simulation 3) and spatially organized populations (Simulation 4). Both led to convergence to Saussurean communication systems, and both seem likely to occur in animal communication. Simulation 4 is particularly important, as it places a lower bound on the level of sophistication an organism would need to evolve this kind of communication. The individuals used in this simulation were extremely simple — just a lookup table without any memory (as opposed to the architecture used in Simulation 3, which had a single bit of memory to keep track of past interactions). While this certainly does not mean that animals exhibiting this kind of communication are this simple, it does mean that a level of complexity higher than that used in these simulations cannot be attributed to such animals solely on the basis of their use of a Saussurean communication system.

References

Ackley, D. and Littman, M., 1991, Interactions between learning and evolution, in: Artificial Life II, C. Langton, C. Taylor, J. Farmer and S. Rasmussen (eds) (Redwood City, CA: Addison-Wesley), pp. 487–509.

Axelrod, R., 1984, The evolution of cooperation (New York: Basic Books).

Axelrod, R., 1987, The evolution of strategies in the iterated prisoner's dilemma, in: Genetic Algorithms and Simulated Annealing, Davis, L. (ed.) Chapter 3, (Los Altos, CA: Morgan Kaufmann Publishers, Inc), pp. 32–41.

Belew, R., 1990, Evolution, learning, and culture: Computational metaphors for adaptive algorithms. *Complex Systems* 4, 11–49.

Cheney, D. and Seyfarth, R., 1990, How monkeys see the world: Inside the mind of another species (Chicago: Univ. of Chicago Press).

Collins, R. and Jefferson, D., 1991, Antfarm: Towards simulated evolution, in: Artificial Life II, C. Langton, C. Taylor, J. Farmer and S. Rasmussen (eds) (Redwood City, CA: Addison-Wesley), pp. 579–601.

Dawkins, R., (1989), The selfish gene (Oxford: Oxford University Press).

Grim, P., 1996, Spatialization and generosity in the stochastic Prisoner's Dilemma. *BioSystems* 37, 3–17.

Hamilton, W., 1964, The genetical evolution of social behaviour (i and ii). *J. Theor. Biol.* 7, 1–52.

Hinton, G. and Nowlan, S., 1987, How learning can guide evolution. *Complex Systems* 1, 495–502.

Holland, J., 1975, Adaptation in natural and artificial systems. (Ann Arbor, MI: The Univ. of Michigan Press).

Hurford, J., 1989, Biological evolution of the saussurean sign as a component of the language acquisition device. *Lingua* 77, 187–222.

Hutchins, E. and Hazelhurst, B., 1991, Learning in the cultural process, in: Artificial Life II, C. Langton, C. Taylor, J. Farmer and S. Rasmussen (eds) (Redwood City, CA: Addison-Wesley), pp. 689–706.

Lindgren, K., 1991, Evolutionary phenomena in simple dynamics, in: Artificial Life II, C. Langton, C. Taylor, J. Farmer and S. Rasmussen (eds) (Redwood City, CA: Addison-Wesley), pp. 295–311.

Nolfi, S., Elman, J. and Parisi, D., 1990, July, Learning and evolution in neural networks. Technical Report 9019, CRL, University of California, San Diego.

Oliphant, M., 1994, Evolving cooperation in the non-iterated prisoner's dilemma: The importance of spatial organization. In Artificial Life IV, R. Brooks and P. Mues (eds) (Cambridge, MA: MIT Press), pp. 349–352.

Parisi, D., Nolfi, S. and Cecconi, F., 1991, June, Learning, behavior, and evolution. Technical Report PCIA-91-14, CNR, Rome.

Saussure, F. de, 1959, Course in general linguistics (New York: McGraw-Hill).

Werner, G. and Dyer, M., 1991, Evolution of communication in artificial organisms, in: Artificial Life II, C. Langton, C. Taylor, J. Farmer and S. Rasmussen (eds) (Redwood City, CA: Addison-Wesley), pp. 659–687.

Yanco, H. and Stein, L., 1992, An adaptive communication protocol for cooperating mobile robots. In From animals to animals, J.A. Meyer, H.L. Roitblat and S.W. Wilson (eds) (Cambridge, MA: MIT Press) pp. 478–485.